

INTERNET DOCUMENT INFORMATION FORM

A . Report Title: Remote Observing with the Keck Telescope
Using the ACTS Satellite

B. DATE Report Downloaded From the Internet 9/22/98

**C. Report's Point of Contact: (Name, Organization, Address,
Office Symbol, & Ph #):** Nasa Lewis Research Center
21000 Brookpark Road
Cleveland, OH 44135-3127
ATTN: Doug Hoder (216) 433-8705

D. Currently Applicable Classification Level: Unclassified

E. Distribution Statement A: Approved for Public Release

**F. The foregoing information was compiled and provided by:
DTIC-OCA, Initials:** VM **Preparation Date:** 9/23/98

The foregoing information should exactly correspond to the Title, Report Number, and the Date on the accompanying report document. If there are mismatches, or other questions, contact the above OCA Representative for resolution.

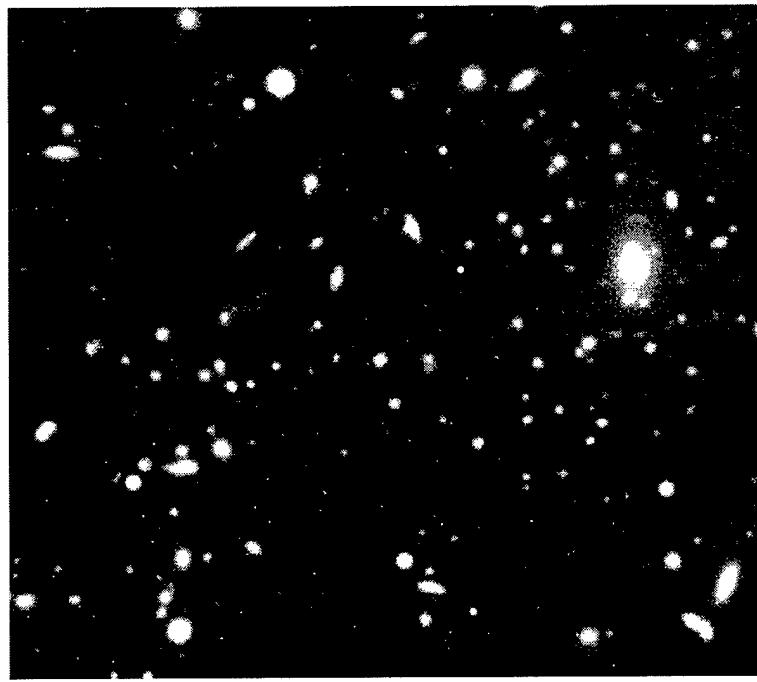
Remote Observing with the Keck Telescope Using the ACTS Satellite

Final Report for NASA Grant #BK-509-20-42-00-00

Judy Cohen¹, Patrick Shopbell²,
California Institute of Technology, Pasadena, California

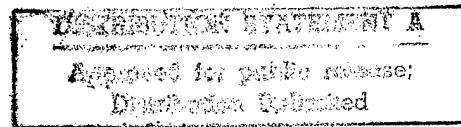
Larry Bergman³
Jet Propulsion Laboratory, Pasadena, California

March 1, 1998



One of the first images obtained via the ACTS/Keck remote observing project. This cluster of galaxies, known as Abell 963, is at a distance of approximately 3 billion light years from Earth.

¹jlco@astro.caltech.edu
²pls@astro.caltech.edu
³larry@jplopto.jpl.nasa.gov



I 98-12-2578

19980925 041

Abstract

As a technical demonstration project for the NASA Advanced Communications Technology Satellite (ACTS), we have implemented remote observing on the 10-meter Keck II telescope on Mauna Kea in Hawaii from the California Institute of Technology campus in Pasadena. The data connection consists of optical fiber networks in Hawaii and California, connecting the end-points to high data rate (HDR) ACTS satellite antennae at JPL in Pasadena and at the Tripler Army Medical Center in Honolulu. The terrestrial fiber networks run the asynchronous transfer mode (ATM) protocol at DS-3 (45 Mbit/sec) speeds, providing ample bandwidth to enable remote observing with a software environment identical to that used for on-site observing in Hawaii.

This experiment has explored the data requirements of remote observing with a modern research telescope and large-format detector arrays. While the maximum burst data rates are lower than those required for many other applications (e.g., HDTV), the network reliability and data integrity requirements are critical. As we show in this report, the former issue particularly may be the greatest challenge for satellite networks for this class of application. We have also experimented with the portability of standard TCP/IP applications to satellite networks, demonstrating the need for alternative TCP congestion algorithms and minimization of bit error rates (BER).

Reliability issues aside, we have demonstrated that true remote observing over high-speed networks provides several important advantages over standard observing paradigms. Technical advantages of the high-speed network access include more rapid download of data to a user's home institution and the opportunity for alternative communication facilities between members of an observing team, such as audio- and videoconferencing. Scientific benefits include involving more members of observing teams while decreasing expenses, enhancing real-time data analysis of observations by persons not subject to altitude-related conditions, and providing facilities, expertise, and personnel not normally available at the observing site. Although the current bandwidth of the public Internet is insufficient for true remote observing between Hawaii and the mainland U.S., we nevertheless anticipate a growing role for remote observing techniques, particularly as high-speed terrestrial networking paradigms, such as ATM, become more commonly available.

1 OVERVIEW

1.1 Remote Observing

Remote use of astronomical telescopes has been a topic of interest for many years, even before space-based observing platforms (e.g., IUE) began to demonstrate total remote operation out of sheer necessity. Initially, a number of ground-based radio and optical telescopes (e.g., the WIYN Telescope [8]) introduced queue-based scheduling, a mixture of remote and interactive observing modes. Only very recently are optical telescopes beginning to realize the benefits of true remote observing: for example, observations with modest size detectors at Apache Point Observatory are being carried out remotely using the Internet [11]. For this project, we have established remote interactive observing capabilities for Keck Observatory on Mauna Kea for observers at Caltech, in Pasadena, California. The recently commissioned twin 10-meter Keck Telescopes are the largest optical/infrared telescopes in the world and thereby typify the data and network requirements of a modern observatory. In undertaking this project, we were motivated by several operational and scientific advantages that remote observing would offer.

One primary concern is the high altitude of the Keck Observatory. At an elevation of 13,600 feet, the summit of Mauna Kea is a demanding location for both mental and physical exertion. In spite of the requirement that all astronomers spend a night at Hale Pohaku (altitude \approx 9000 feet) for acclimatization before proceeding to the summit for a night of observing, about 15% of the people who do not observe often at Mauna Kea become sufficiently ill during the course of a 3 night run that they have to leave the summit for at least 12 hours. Approximately 75% of the people coming to the summit to observe for a full night experience some discomfort such as a mild headache, and almost all experience some loss of judgment, irritability, etc. Remote observing provides an environment for all observers that is free of these difficulties, and also provides an opportunity for people who cannot tolerate high altitudes (e.g., pregnant women, those with heart conditions, etc.) to observe with the Keck Telescopes.

Another logistic motivation for remote observing involves monetary issues common to large telescopes located at sites distant from the home institutions: In general, the larger the telescope, the more heavily over-subscribed it is. Runs are therefore often only 1 or 2 nights in duration. Since 2–3 observers come to each run, this means substantial sums of money are spent on travel and related expenses. The additional night of acclimatization for high-altitude sites such as Mauna Kea increases the cost further. Finally, the salary cost for “wasted time” during these runs is quite large. An excellent example of the potential savings is that of the European Southern Observatory (ESO): There are 19 telescopes near the Atacama desert in Chile, including two 3.6-meter telescopes and a set of four 8-meter giants currently under construction (the Very Large Telescope, or VLT). The observing site is widely regarded as one of the very best in the world, yet it is half-way around the globe from most of its large European user population. Understandably, remote observing is gaining popularity among European astronomers [12].

Remote diagnosis of hardware and software problems also becomes more feasible with an operational remote observing system. In the case of Keck Observatory, the teams that built the instrument hardware and software are located at Caltech or a campus of the University of California. Just their presence in the same buildings as the remote observers can be extremely helpful when problems arise in the operation of the telescope. The establishment of remote observing from California implies the presence of a network connection, which can allow engineers and programmers to analyze the remote systems essentially instantaneously. Again, both travel and time are saved, and effective help from highly skilled and experienced people in California can be obtained quickly when necessary.

In addition to these operational advantages, there are strong scientific advantages to remote observing as well. With remote observing, every member of a large collaboration can participate in obtaining the data. It is possible for one part of the team to concentrate on obtaining the observations, while other team members can be analyzing the scientific results from the last observation, checking the instrumental performance to make sure everything is working correctly (particularly the detector), or browsing the literature or catalogs of objects as necessary to prepare for the next set of observations. The inclusion of students in the observing session becomes much easier, cheaper, and more routine when no travel is required, i.e., they don't need to miss classes. The facilities available at remote observing sites (e.g., Caltech) usually far exceed those available at the observatory site, whether it be computer hardware, office and library supplies, or a pizza delivery. Recall also that the remote site may be located such that the night hours of observing overlap, or even coincide with standard business hours at the remote site.

1.2 Network Requirements

Given these strong logistic, financial, and scientific motivations for remote observing, we may then explore the network requirements to make remote observing feasible. The primary issue involves the large size of modern astronomical images. The optical instruments in use at the Keck Telescope have frames that are currently 8 Megabytes in size, soon to become a factor of 4 larger. Although actual integration times depend on the scientific program, and range from less than a second to an hour, the quality of ground based optical and infrared astronomy observations is very sensitive to weather conditions, including clouds and atmospheric turbulence. Hence, even though observing sessions are planned in detail in advance, careful "quick look" analysis of each image is important in defining what to do next, how long the next exposure should be, whether to switch to brighter objects due to poor sky conditions, how to modify the program to cope with unexpected failures of non-critical telescope or instrument components, etc. An operating mode such as this clearly requires a means of viewing or retrieving the images at the remote observing site with a minimal amount of delay. The public Internet connection between Hawaii and California was in 1995 (and still is today) insufficient for these purposes.

Beyond the network requirements for rapid image data transfer, telescope instrument control software generally employs minimal bandwidth. Due to stringent requirements on robustness and ease of use, long software development times, and the wide variety of astronomical instrument characteristics, instrument software interfaces are rarely more complex than a single interactive window. In some cases that window may even be the user's web browser, as recently groups have been experimenting with front-end instrument control interfaces based on Sun Microsystems' Java language (e.g., [9]).

Finally, previous remote observing projects have demonstrated the need to maintain a strong communications link between the remote astronomer and any on-site technical or operations personnel. Not only does the astronomer require adequate communications to direct the course of the observing run, but the on-site staff also value the contact and stimulation that interaction with scientists brings. There are a range of solutions for this issue, spanning a range of bandwidth requirements, from a simple text-based "chat" window, to full audio/video-conferencing systems. Regardless, it is crucial that the communications link not interfere with the accurate transmission of scientific data.

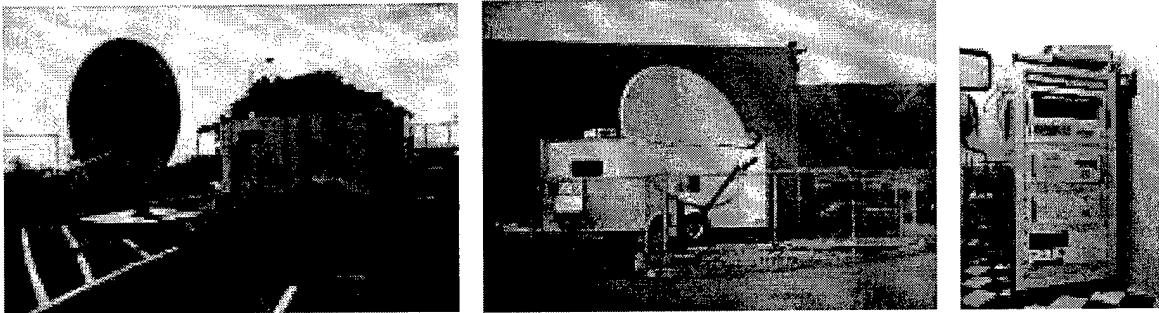


Figure 1: The ACTS high data rate (HDR) ground stations: *a.* The HDR at Tripler Army Medical Center in Honolulu, Hawaii. Visible are the 12-foot antenna and associated equipment trailer, connected by a dry waveguide. Note the very low elevation of the antenna, due to the extreme western longitude of Hawaii. This low elevation exacerbated rain fade problems. *b.* The HDR at JPL in Pasadena, California. *c.* The control unit inside each HDR equipment trailer. From top to bottom, the rack contains a real-time Unix control system with SONET I/O boards, a burst modem, and a high-output transmitter.

1.3 The ACTS Satellite

In light of these requirements, we submitted a proposal to NASA as part of its Advanced Communications Technology Satellite (ACTS) Gigabit Satellite Network (GSN) testbed program. We received funding to establish a high-speed ATM network running from the domes of the 10-meter Keck Telescopes on the summit of Mauna Kea in Hawaii to the Caltech campus in Pasadena, California, using the ACTS satellite as the network link across the Pacific Ocean.

The main deterrent to the implementation of remote observing has always been the problem of obtaining an affordable and reliable connection with adequate bandwidth. NASA's Advanced Communications Technology Satellite was built as a prototype system to explore new modes of high speed transmission for digital data. It provides this capability at rates reaching up to OC-12 (622 Mbit/sec) via advanced on-board switching and multiple dynamically hopping spot beam antennae for selected areas of the United States, including Pasadena and Hawaii, although the steerable antenna used to reach sites not in the continental U.S. is only capable of OC-3 (155 Mbit/sec) speed. The 20–30 GHz frequency band has been employed for the first time by a communication satellite, with extensive rain fade compensation.

ACTS was launched on September 12, 1993 by Space Shuttle Discovery and now occupies a geostationary orbit at 100° west longitude. It has survived almost twice as long as its planned mission duration of two years, but is now nearing the end of its lifetime, which is limited by the fuel resources required to maintain its stationary position. (Current plans involving steerable ground stations may be implemented to extend the usable lifetime of the satellite even further.) The ACTS program is administered by NASA's Lewis Research Center (LeRC) in Cleveland, Ohio.

Bolt, Beranek, and Newman, Inc. (BBN) designed, built, and maintains the high data rate (HDR) ground stations that provide a gateway between ACTS and ground-based fiber optic networks and supercomputer interfaces. Five of the semi-portable HDR terminals have been built; they are allocated to the various ACTS experiments for predetermined lengths of time, then moved to another location. (For more information on the ACTS satellite and program, see the Gigabit Satellite Network web page¹.) Each HDR ground station includes a 12-foot dish permanently pointed at the satellite and an equipment trailer containing a real-time Unix control system with SONET I/O boards, burst modem, and high-output transmitter (see Figure 1). Due to the ex-

¹<http://mrpink.lerc.nasa.gov/gsnhome.html>

perimental nature of these ground stations, the often harsh environmental conditions, and the inherent complexity of high-speed communications equipment, the HDR stations have proved to be the weakest link in our network.

This network has been used to support remote observing, remote diagnosis of problems, remote software development, and other related tasks. In the sections that follow, we will outline the network architecture and topology, characterize the performance of the network, demonstrate remote operation of a specific instrument on the Keck II Telescope, and suggest future directions for remote observing with high-speed networks. We will close with a summary of the benefits and difficulties which we have encountered during the course of this ACTS demonstration project.

1.4 Participants

Primary participants in this project included the California Institute of Technology, the Jet Propulsion Laboratory (JPL), and Keck Observatory. Other important contributions have been made by the ACTS program, Tripler Army Medical Center (Honolulu), the Pacific Space Center (PacSpace, Honolulu), Pacific Bell, and GTE Hawaiian Telephone. Many software and hardware vendors assisted in debugging various aspects of the network, including Sun Microsystems, Fore Systems, Newbridge Networks, and the Mitre Corporation.

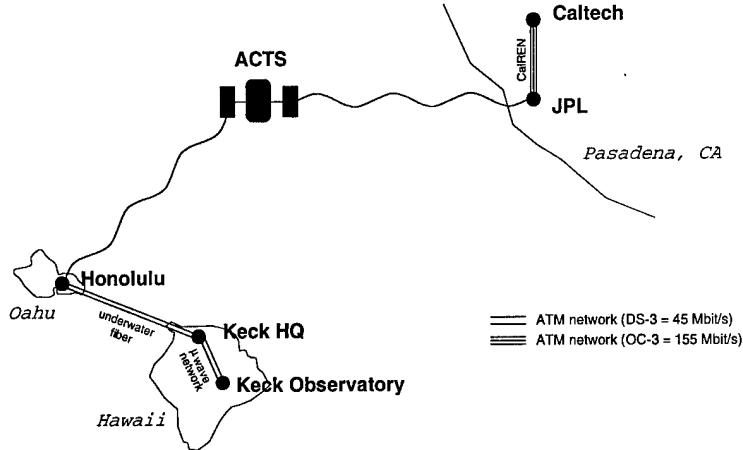


Figure 2: A schematic of the initial high-speed ACTS network used for the Keck remote observing project. The entire ground-based network now consists of optical fiber.

2 NETWORK ARCHITECTURE

We now describe the topology of our remote observing network, including the hardware infrastructure, software protocols, and user-level application interface.

2.1 Physical Layer

Following 6 months of installation of optical fiber, satellite stations, and microwave antennae, our network began end-to-end testing in February of 1996. The network consists of three major segments: the ground network in California, the satellite link across the Pacific Ocean, and the ground network in Hawaii (see Figure 2).

The ground network in California connects Caltech with JPL, the site of the satellite ground station. This portion of the network was established as part of Pacific Bell's extant fiber optic network. Due to the integrated nature of Caltech and JPL, the only infrastructure required to establish this physical connection was the installation of a fiber optic line from the Caltech backbone to the remote observing room in the astronomy building. Existing available bandwidth between Caltech and JPL well exceeded our requirements. This segment of the network has been extremely stable, remaining reliable and unchanged for the duration of our experiment.

The ground network in Hawaii has been somewhat more complex in its evolution, primarily due to the relative inexperience of GTE Hawaiian Telephone, as compared to PacBell in California, and a lack of prior infrastructure in Hawaii. The first segment of the Hawaii network consisted of undersea optical fiber connecting the satellite ground station in Honolulu to the GTE Hawaiian Telephone office on the big island of Hawaii. Although already in existence, this fiber had been installed less than a year before our project began. The next segment of the network utilized microwave antennae to reach across the big island of Hawaii to Hale Pohaku, at the 9,000-foot level on Mauna Kea (see Figure 3a.). At that time, fiber optic cable had not yet been installed in these relatively remote areas. As we shall show, the introduction of the higher bit error rates (BER) of this non-fiber segment produced noticeable instability in the end-to-end network. Fortunately, in January of 1997 this portion of the ground network in Hawaii was upgraded to optical fiber. The improved performance for high-speed data transfers of the final all-fiber network was immediately

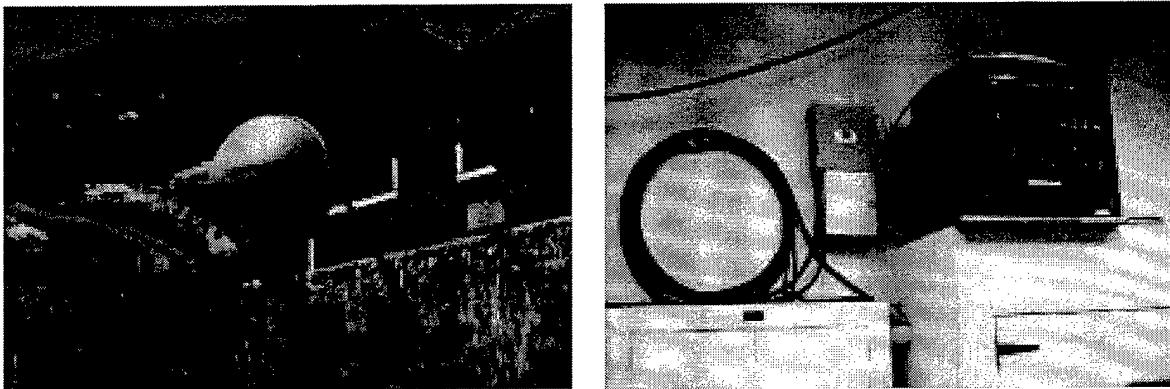


Figure 3: Hawaiian network infrastructure: *a.* A microwave relay station at Hale Pohaku, at the 9,000-foot level on Mauna Kea. Before the completion of a fiber network up the mountain in January of 1997, several dishes such as this one relayed data between Hale Pohaku and the base of the mountain. *b.* The ATM switch installed on the summit of Mauna Kea, in the Keck Telescope building. The switch is a Newbridge 36150 MainStreet ATMnet, provided by GTE Hawaiian Telephone. Two single-mode DS-3 cards are installed in the left end of the switch.

apparent. The final segment of the Hawaii network, from Hale Pohaku to the telescope dome on the summit of Mauna Kea, employed pre-existing fiber optic cable. Figure 4 illustrates the final network configurations in Hawaii and California.

2.2 Network Protocol Layer

The physical-level protocol used by this network is the standard Synchronous Optical Network (SONET) optical data transmission protocol. This is the level at which the ACTS satellite operates, i.e., it knows nothing of protocols above the raw SONET data stream.

In order to support standard higher-level (IP) networking protocols, we installed an Asynchronous Transfer Mode (ATM) network on top of the SONET layer. ATM is a packet-switched protocol, similar to frame relay, which is capable of bandwidths exceeding OC-48 (2 Gbit/sec) [2, 10]. Data is transferred in 53-byte "cells", each containing a 5 byte header and 48 bytes of payload (see Figure 5). The transfer of cells is performed by hardware switches, which have been installed throughout the network in California and Hawaii. The ground network in California includes three ATM switches running at OC-3 (155 Mbit/sec) speeds: one at each end-point (Caltech and JPL), and an intermediate switch belonging to the CalREN (California Research and Education Network) project of Pacific Bell. The ground network in Hawaii includes several ATM switches: one at each end-point (Honolulu, the Keck Telescope dome, and the Keck Headquarters in Waimea; see Figure 3*b.*), and a number of intermediate switches belonging to GTE Hawaiian Telephone. Because of the initial use of microwave antennae in Hawaii, this portion of the network, and therefore the end-to-end network, was limited to DS-3 (45 Mbit/sec) speeds.

The ATM switches provide a point-to-point network connecting the computer in the remote observing room at Caltech with the instrument control computer at the Keck Telescope in Hawaii. We have established Permanent Virtual Circuits (PVCs) in each of the switches, which direct the cells between the end-points of the network. Several vendors have supplied the ATM switches for this network, including FORE Systems, Newbridge Networks, and SynOptics. Fore ATM Network Interface Cards (NICs) are used to interface the Sun SPARCstation 20/51 workstations at each end of the network. This mixed vendor environment has been a stringent test of the compatibility

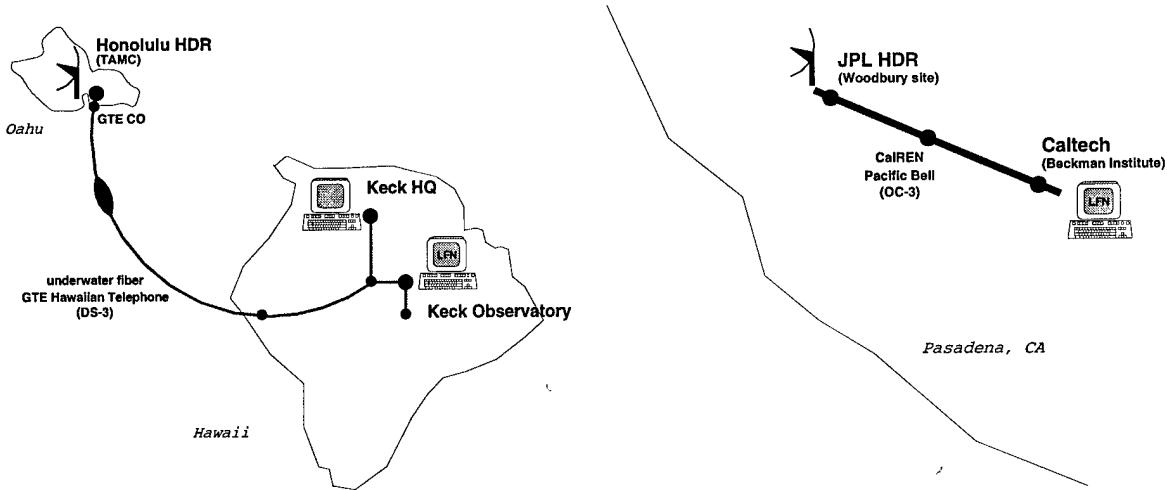


Figure 4: Schematics of the final terrestrial networks in Hawaii and California, as used for the Keck remote observing project. The Hawaii network connects the Keck Observatory to the ACTS HDR station in Honolulu, at DS-3 (45 Mbit/sec) speeds. The California network connects the Caltech campus to the ACTS HDR at JPL, at OC-3 (155 Mbit/sec) speeds. Black dots denote ATM switches.

among vendors in the relatively new ATM environment. Although we have encountered several problems associated with interoperability issues, none have been extremely serious, and the ATM vendors and telephone companies have been extremely helpful in attempting to diagnose and solve ATM-level problems. We have also witnessed the increasing popularity of ATM even in the mere 2-year lifetime of this project, and expect to see more widespread use of this protocol at the WAN and enterprise network levels. (For more information on ATM, see the Cell Relay web site².)

2.3 User Protocol Layer

ATM is not a "reliable" protocol in the networking sense - bit errors and congestion may cause cells to be dropped or lost in transit; no attempt is made to verify delivery. In order to facilitate reliable data transfer for scientific applications, as well as to allow the use of the wealth of software tools already available, we are running the standard IP protocols over ATM. We are using a pseudo-standard implementation known as "Classical IP", which defines a relationship between standard IP "dot" addressing and ATM PVCs, as well as data packet segmentation and re-assembly algorithms for converting between IP packets and ATM cells [1] (i.e., AAL-5). Since the ATM switches know nothing of upper-level protocols, the choice of IP as a user protocol affects merely the two end-point systems. Those workstations both employ FORE SBA-200 ATM interface cards to perform IP packet segmentation and re-assembly in hardware. These interface cards run at speeds up to OC-3 (155 Mbit/sec).

Over the Classical IP level, we run the standard TCP/IP and UDP/IP protocol suites. This enables the use of all the standard network-based applications that are in widespread use on the Internet. Common tools such as *ftp*, *telnet*, and the *X Window System* are part of every observing run, as are additional special-purpose applications, such as an audio conferencing tool (*rat*) and a shared whiteboard tool (*wb*).

The most important impact of a satellite component on a high-speed network is the relatively

²<http://cell-relay.indiana.edu/cell-relay/>

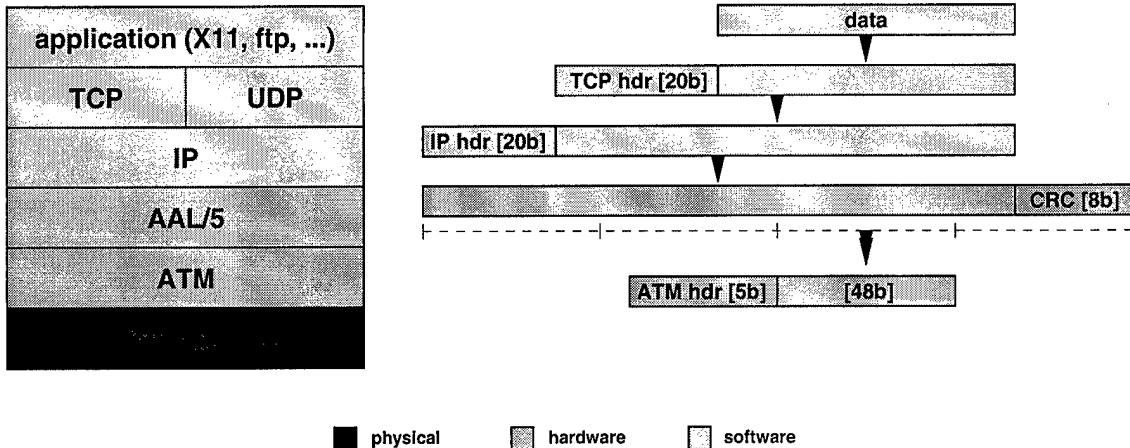


Figure 5: The protocol stack for our ATM-based TCP/IP network.

large delay introduced by the round-trip signal travel time to the satellite. In our network, this travel time is approximately 0.55 seconds, which corresponds to over 3 Mbytes of data at DS-3 speeds (45 Mbit/sec). The problem has to do with the connection-oriented nature of TCP/IP: TCP sends a very specific amount of data, known as a "window", after which time it expects an acknowledgment from the other end of the connection. This is the manner in which TCP/IP is able to implement a "reliable" connection. However, this window size is often very small; the default value for workstations running the SunOS 4.1.4 operating system is 4 Kbytes. If one were to use this system on our satellite network in its default configuration, a window of data would be sent in 0.0007 seconds, after which the sending system would be forced to wait 0.549 seconds for an acknowledgment. In other words, the system would be running at 0.1% efficiency, and the net throughput would reflect this: initial tests of our system under such conditions showed bandwidths of 0.1–0.2 Mbit/sec.

Fortunately, this problem is well-known in the high-speed networking community. Networks such as ours are known as "long fat networks" (LFN). The figure of merit for such networks is the window size, or the bandwidth-delay product:

$$\text{TCP window size} = \text{bandwidth} * \text{delay time} \quad (1)$$

(See Figure 6.) There is an Internet Request For Comment (RFC) document on this subject, RFC 1323 [5], and the problem has been discussed extensively. Many current operating systems support the RFC 1323 extensions, and provide options to increase the TCP window size to values in excess of 10 Mbytes. In the case of the SunOS operating system (to which we are constrained by legacy control software at Keck), we obtained the TCP-LFN package from Sun Consulting, which also purports to support the RFC 1323 extensions for long fat networks.

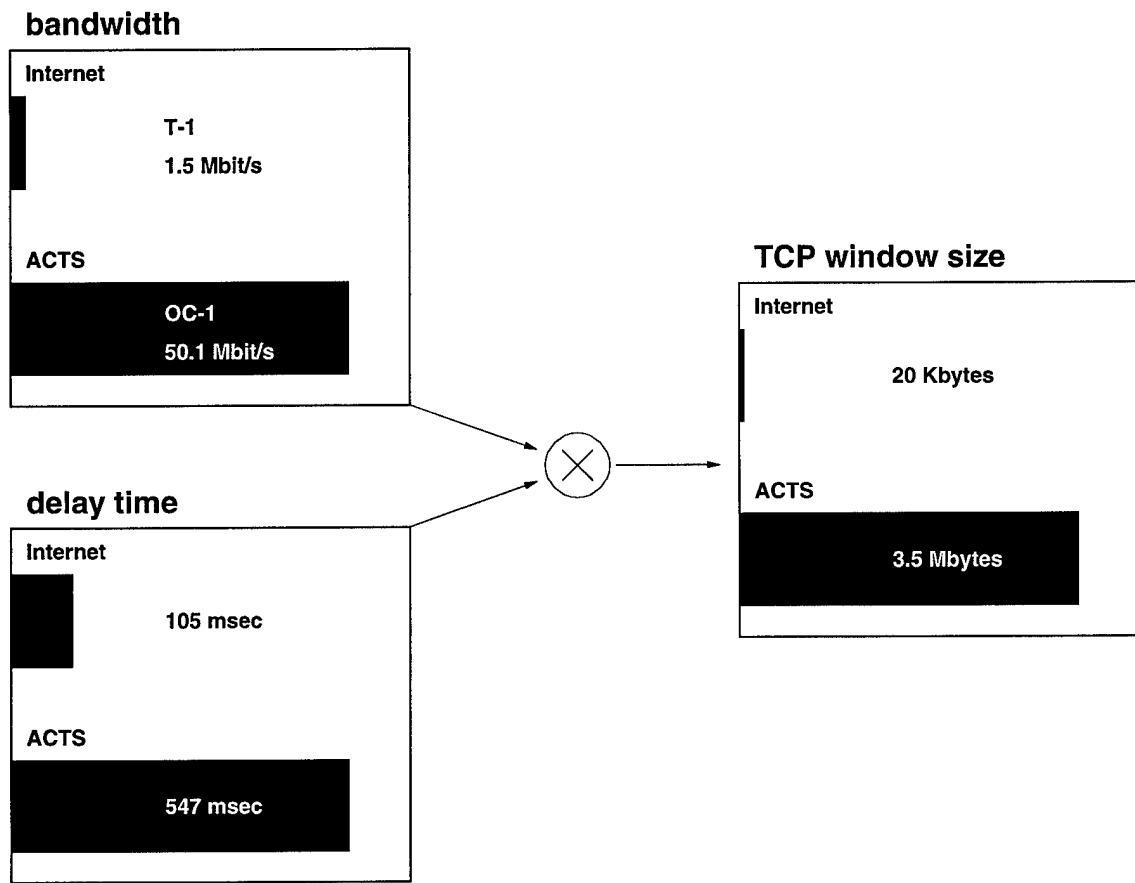


Figure 6: In order to increase the TCP acknowledgment window to values greater than 64 Kbytes, extended TCP windows (i.e., RFC 1323) must be supported at the operating system level. This support is virtually required for satellite networking, especially in high-bandwidth applications. For the Keck/ACTS network, the optimal extended TCP/IP window size is approximately 3.5 Mbytes.

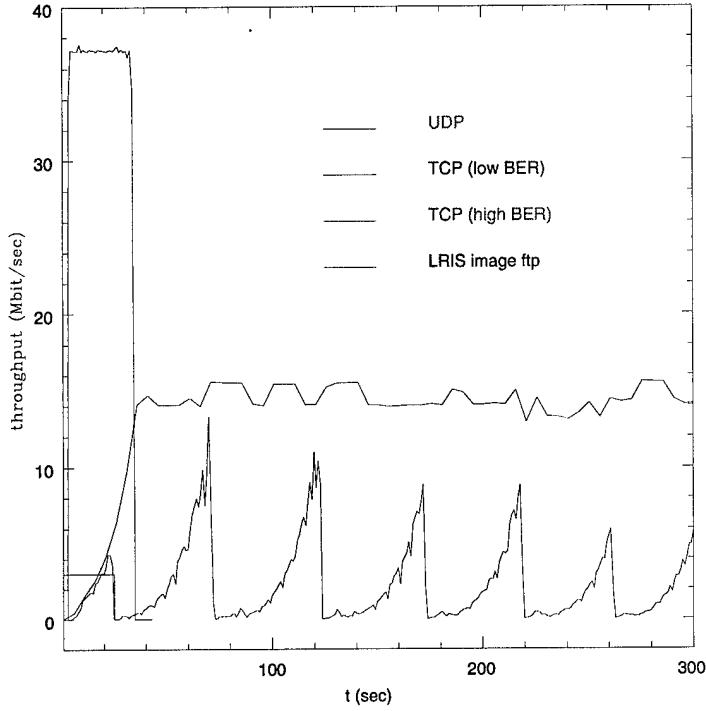


Figure 7: Bandwidth test results between Keck Observatory and the Caltech campus in Pasadena, California, over the ACTS satellite network. Measurements of UDP and TCP transfers of a fixed-length data stream are shown. TCP exhibits a remarkable dependence on the bit error rate, as is demonstrated by measurements before and after the conversion of microwave antennae to fiber optic cable in Hawaii. Also evidenced is the need to select an adequate TCP window size for satellite networks. Because of limitations in the SunOS kernel, we have been limited to a TCP window size of 1 Mbyte, approximately one-third of the preferred value for this network.

3 NETWORK PERFORMANCE

The performance of the network was gauged using standard network tools, the primary one being the freeware `ttcp` utility. This utility measures the bandwidth of a network connection via memory-to-memory host data transfers, producing measures that are independent of disk speeds. The resulting statistic is a product of the processor speeds of the end-point host computers, the intrinsic speed of the underlying network fabric, and the efficiency of the lower-level protocols in terms of the amount of packaging overhead.

The issue of end-point host processor speed was known in the beginning of the project, and we obtained the fastest machines then available which were also compatible with the Keck Observatory control software. These SPARCstation 20/51 workstations were also equipped with 1 Mbyte of level 2 processor cache to increase network throughput.

The second issue of importance in assessing the performance of the network concerns the intrinsic speed of the network fabric itself. As has been discussed, the California ATM network between Caltech and JPL was configured to run at speeds up to OC-3 (155 Mbit/sec). We confirmed this number through simple tests between our Caltech end-point and a JPL Cray system

at the HDR site: with little tuning we were able to measure effective TCP and UDP bandwidths in excess of 85 Mbit/sec. Similarly, the ACTS satellite connection between California and Hawaii is configured to run at OC-3 speeds. (Although ACTS is capable of OC-12 [622 Mbit/sec] communication, the steerable antenna which reaches Hawaii is capable of only OC-3 speeds). Again, this speed was measured using our end-point host and the JPL Cray, with ACTS placed in a "bent pipe" configuration to connect the two. In contrast, the Hawaii ATM network was configured to run at only DS-3 (45 Mbit/sec) speeds. Although originally the network was intended to run at these speeds only while the microwave antennae were needed on the big island of Hawaii, a lack of OC-3 interface cards for GTE Hawaiian Telephone's ATM switches prevented us from attempting to increase the speed of the Hawaii network during the later stages of the experiment. This limitation set the maximum speed for our network at 45 Mbit/sec (DS-3).

Finally, since our performance measurements are computed based on transmission speeds of actual user data, the results also reflect the amount of packaging overhead in the lower-level protocols. In the case of TCP packets, this overhead includes TCP and IP headers (20 bytes each), an ATM CRC (8 bytes), and an ATM header in each cell (5 bytes). (See Figure 5.) This issue was confronted by adjusting a number of TCP and IP parameters to minimize the fractional overhead. First, we used a large TCP packet size of 65536 bytes for all testing. Unfortunately, this may give slightly skewed results, as it is difficult to modify the packet size used by the end-point systems in non-testing situations: the systems' network drivers will adjust the TCP packet size dynamically in an attempt to optimize throughput. However, this parameter is not extremely important, as TCP packets are broken up into smaller segments for transmission. The second modification we made was to raise the Maximum Segment Size (MSS) of these segments to 1500 bytes, approximately a factor of 3 above that normally used, and the highest value which is safe to assume routers can handle. Thus, each 1480 bytes of user data will be accompanied by a 20-byte TCP header for that segment. Finally, the Maximum Transmission Unit (MTU) for IP was increased to 9180 bytes. In the case of TCP, any value in excess of the MSS (plus 20 bytes of IP header) is sufficient to ensure that each TCP packet is transmitted within a single IP packet. In the case of UDP, this value limits the quantity of UDP data which may be transmitted in a single IP packet to 9160 bytes.

Given these values, we may then calculate the TCP/IP data transmission efficiency for our network. The number of 53-byte ATM cells required to transmit a single TCP segment is given by:

$$n = (\text{data bytes} + \text{TCP hdr} + \text{IP hdr} + \text{AAL/5 CRC})/\text{ATM data size} \quad (2)$$

$$= (1480 + 20 + 20 + 8)/48 \quad (3)$$

$$= 32 \text{ cells.} \quad (4)$$

Therefore, the efficiency (ratio of data bytes to transmitted bytes) is:

$$\epsilon = 1480/(32 \text{ cells} * 53 \text{ bytes/cell}) \quad (5)$$

$$= 87\% \quad (6)$$

This implies a maximum *data* throughput for our network of 87% of DS-3, or approximately 39 Mbit/sec.

The most complex part of optimizing the throughput of our network has involved the TCP-LFN extensions to the SunOS kernel. As mentioned previously, we employed a Sun Consulting special

package to augment the SunOS kernel with extended TCP windows and other capabilities outlined in RFC 1323. Unfortunately, a number of limitations of SunOS 4.1.4 conspire to prohibit one from obtaining extremely large window sizes, regardless of the TCP-LFN software. In our case, the compiled-in kernel limit of 2 Mbytes of Mbuf memory (i.e., IP packet wrappers) turned out to be the major constraint, limiting our window size to no more than 1 Mbyte. This is approximately one-third of the optimal value derived above in Figure /reffig:Ifn. As such, our final tuned network delivered a maximum TCP/IP performance of approximately 15 Mbit/sec, about one-third of the 39 Mbit/sec expected data throughput (Figure 7).

Although perhaps disappointing in a relative sense, this bandwidth is far in excess of T1 Ethernet speed (1.44 Mbit/sec) and allows an 8 Mbyte image to be transferred in approximately 5 seconds. As a further comparison, this bandwidth exceeds by 50% that which is available on the local area Ethernet network at the Keck Telescope itself.

3.1 Reliability Issues

While network performance was perhaps not at the level desired, due to developing infrastructure in Hawaii and idiosyncrasies within the outdated SunOS 4.1.4 operating system, issues of network reliability had far greater impact on our remote observing operation. The experimental and limited nature of the ACTS program has created a number of difficulties which one would almost certainly not face if using a more developed and/or commercial satellite system. For example, we have been forced to await replacement of two transmitters, the operating system, and several other HDR components, due to hardware failures. The Internet connection to the HDR, which is required for operation, has also proved unreliable. Although the ACTS personnel and BBN have been extremely cooperative in restoring service on such occasions, the impact of the reliability issue is that at least one observer must be sent to Hawaii to use the telescope, in case of ACTS-related equipment malfunctions.

The remote nature of most high-quality observing sites exacerbates this problem. BBN maintains only a small field office in Hawaii, making HDR maintenance costly and time-consuming. A truly remote site, which would most benefit from remote observing techniques, also requires the highest degree of robustness from the equipment. In our opinion, the ACTS system is insufficiently robust to provide true remote observing with large (i.e., highly competitive) telescopes, due primarily to its limited scope and experimental nature. However, one of the ACTS Project's primary goals is to stimulate commercial high-speed communications satellite development. These systems may eventually play a role in remote astronomical observing systems.

Another difficulty we have encountered is that the transmitters in the ACTS HDR stations are not designed to run continuously, due to the finite lifetime of certain critical components, but rather must be switched on and off as needed. This method of operation demands human intervention at the beginning and end of every satellite session, a procedure that has been non-trivial to organize and would prove difficult in more remote observatory locations. The Hawaii location itself poses an additional problem due to the large yearly rainfall at the location of the HDR in Honolulu. Because the uplink frequency of 30 GHz at which ACTS operates is highly susceptible to rain fade, we have lost several runs in the past year due to rain in Honolulu. In essence, the use of the ACTS system for remote observing adds a weather constraint such that it must be clear (i.e., not raining) at both of the ground station sites, as well as at the observatory itself! Noting ACTS rain fade compensation capabilities, we suggest that this is another area in which future commercial high-speed communications satellites may provide improvements.

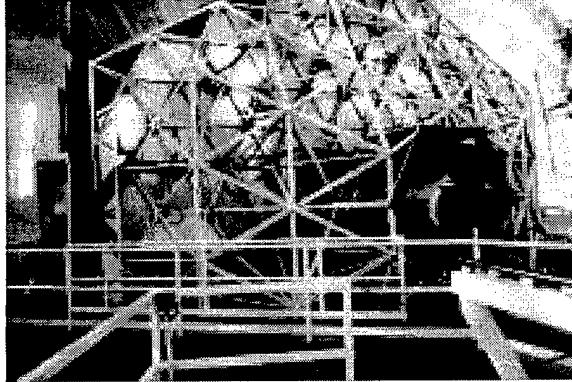


Figure 8: *a.* The domes of the twin 10-meter Keck Telescopes on Mauna Kea, Hawaii. *b.* The back of the primary mirror on the Keck I telescope, illustrating the network of 36 hexagonal segments which make up the mirror. The infrared camera is mounted at the forward Cassegrain focus of the telescope (i.e., in the central hole of the mirror).

4 KECK OBSERVATORY AND THE LRIS INSTRUMENT

The W. M. Keck Observatory, located on the summit of Mauna Kea on the Big Island of Hawaii, consists of twin 10-meter telescopes intended for astronomical observations at optical and infrared wavelengths. Both telescopes employ a revolutionary design, in which each 10-meter mirror consists of 36 2-meter hexagonal segments which are aligned to act as a single large mirror (see Figure 8; [6]). These are the largest astronomical telescopes in the world for use at these wavelengths. The first Keck Telescope has been in routine operation since 1993, and has successfully demonstrated the viability of the multiple-segment mirror design. The second Keck Telescope became operational in October of 1996 and is currently dedicated almost entirely to optical observations.

Throughout the design process of the hardware and software for the Keck Telescopes, the possibility of implementing remote observing, particularly from Waimea, the location of the Keck headquarters in Hawaii, was kept in mind (see Figure 9). The instruments, their motors, and the detectors are operated through workstations that are located in the control room of the Keck Telescope dome. It is not necessary during normal night-time operation to go out to the instrument on the telescope to make any adjustments or changes.

We have concerned ourselves exclusively with enabling remote observing on the Keck II telescope with the Low Resolution Imaging Spectrograph [7] (LRIS; see Figure 10). This is the primary optical instrument at the observatory, and the only instrument capable of obtaining direct images at optical wavelengths. It is used on Keck II almost every night of the year. Although our efforts have concentrated on remote observing with LRIS, we note that all of the instrument interfaces have been engineered in a similar fashion, so our results could be easily extended to other instruments on Keck I or II.

Figure 11 illustrates the organization of the telescope and instrument control hardware at the Keck Observatory. The majority of the complications associated with remote control of the instrument have stemmed from security issues and the desire to not impact normal (i.e., non-remote) observations with the telescope. For example, concerns about the effect of a such a high-speed network, especially in various failure modes, on other computer and network systems at Keck, as we approached this project with our initial inexperience, forced us to isolate remote operations more thoroughly by duplicating the instrument control computer. This machine provides boot infor-

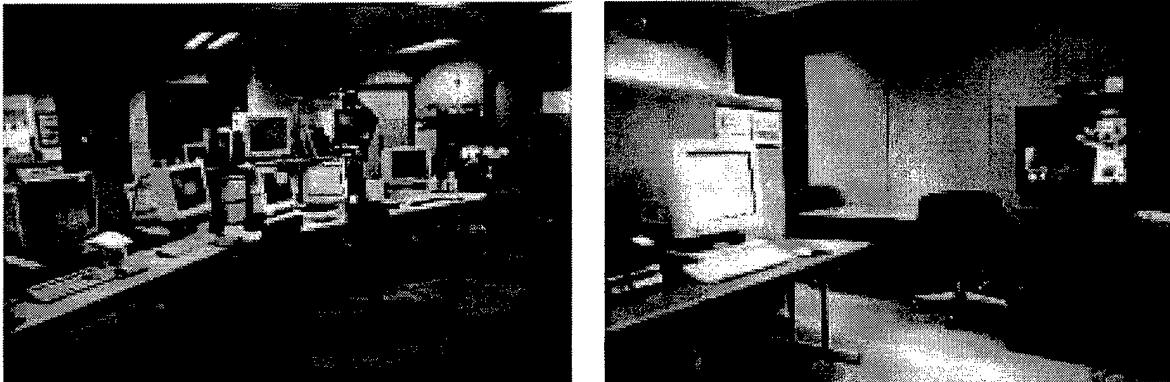


Figure 9: The Keck II telescope control rooms: *a.* in the telescope building itself, on the summit of Mauna Kea, and *b.* in the Keck headquarters building in Waimea, at an altitude of only 2000 feet. The two rooms provide the astronomer with identical control environments for the telescope and its instruments.

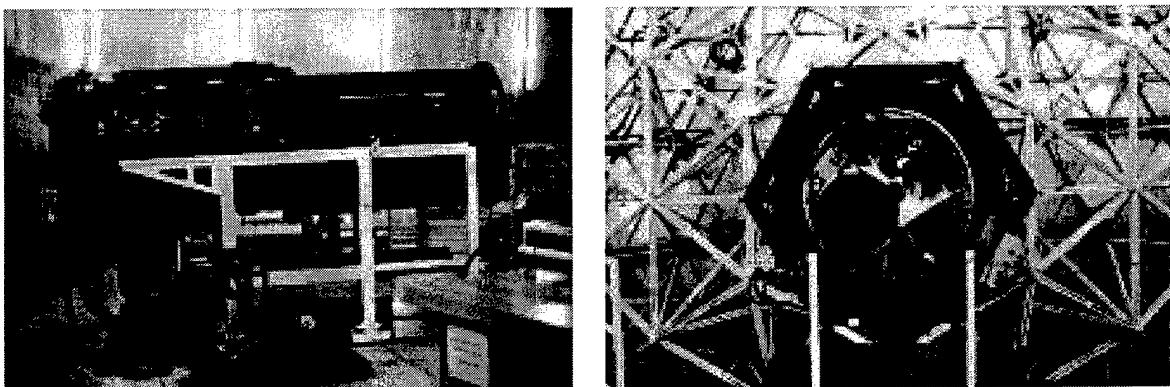


Figure 10: Instruments for the Keck telescope: *a.* the Near-Infrared Camera (NIRC) in a stowed configuration, and *b.* the Low Resolution Imaging Spectrograph (LRIS), installed for operation at the Cassegrain focus of the telescope.

mation for the instrument motor and CCD detector VME crates, as well as the interface between the user and the control systems. The remote portion of the observing system, the workstation at Caltech, is integrated into the system merely as a remote display, using the X Window System protocols. Although this is perhaps not the most efficient method of providing remote operations, it is certainly the most straightforward, especially in a relatively new and frequently changing facility such as Keck.

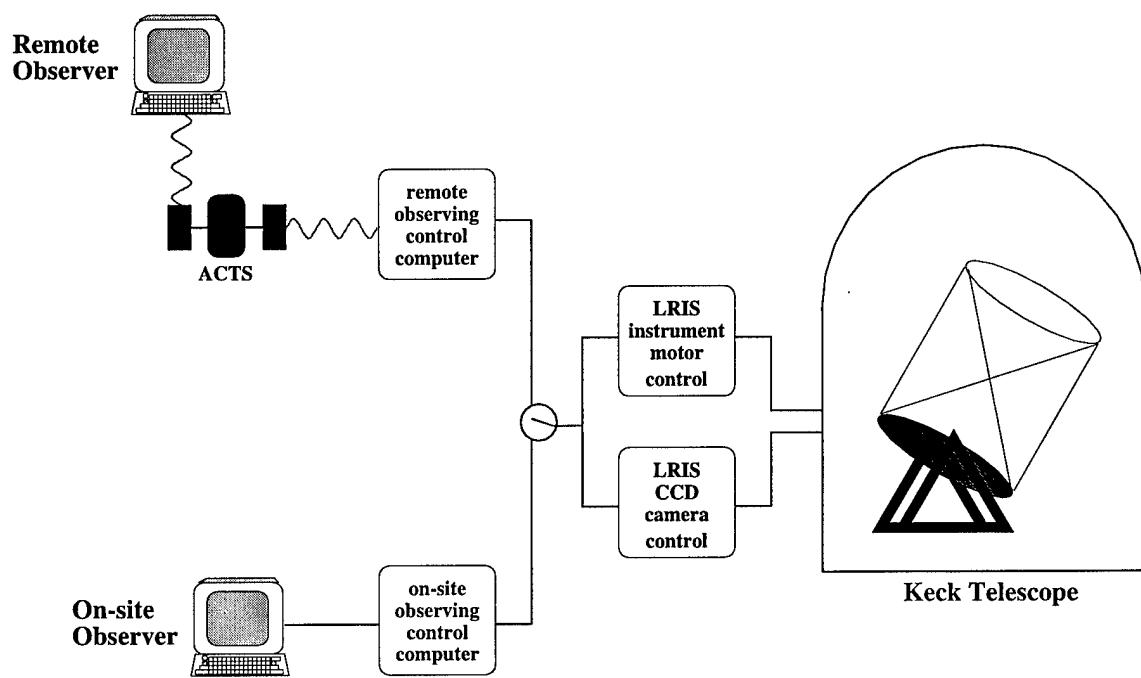


Figure 11: A schematic of the LRIS instrument control system at Keck Observatory. Note that there are two separate instrument control computers: one for standard on-site observations, and another for remote observations over the high-speed network.

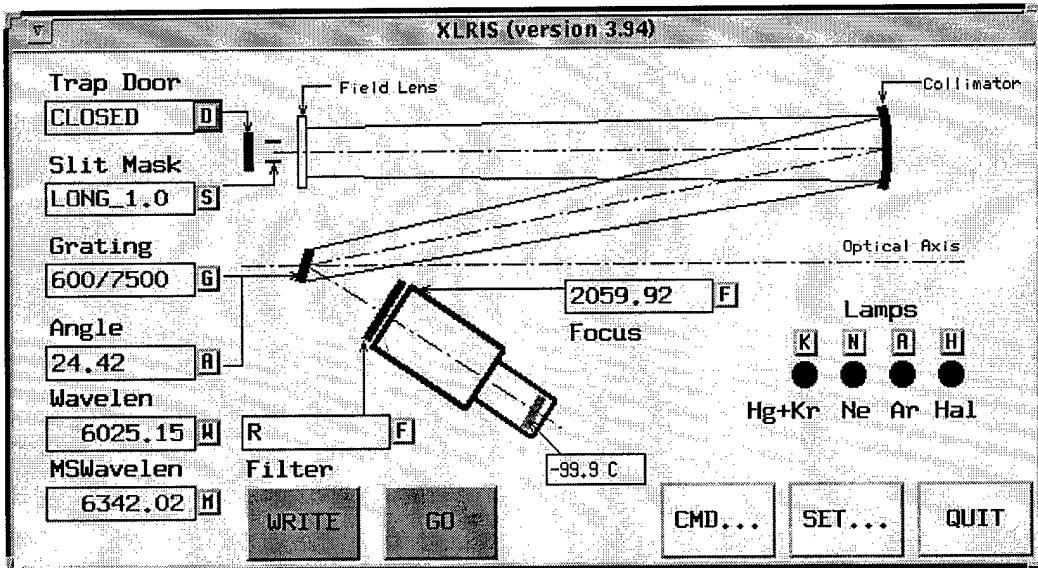


Figure 12: The LRIS instrument control interface at Keck Observatory. This window provides a schematic of the instrument and indicates its current configuration.

5 REMOTE OBSERVING OPERATION

We now outline a night from a typical remote observing run at Caltech, to demonstrate the advantages and problems associated with such operation.

ACTS scheduling For the testbed ACTS experiments such as ours, the satellite must be scheduled 1-2 weeks ahead of time. This is not a problem in most cases, since observing schedules at Keck are established 6 months at a time. The primary difficulty in scheduling ACTS sessions lies in the 5-hour time difference (6 hours during daylight savings time) between Hawaii (HST) and the East coast (EST). Since the satellite experiences higher demand during daylight hours, it is often difficult to run an observing session remotely during the second half of the Hawaiian night. We therefore often restrict remote runs to the first half of the night. But even this is not a critical problem, as the University of California actually allocates its Keck time in half nights, to provide more astronomers with an opportunity to use the telescopes. And in many full-night runs, the first half of an observing night is the most complex and demanding, so eavesdropping and collaborative use of the remote capabilities are very useful in that capacity.

ACTS setup When the time comes for the ACTS session to begin, operators at JPL and at the Tripler Army Medical Center in Honolulu will turn on the transmitters at the HDR ground stations. We are fortunate that both HDR units are located at facilities which are in continuous 24-hour operation every day, so that staff are available to turn on/off the ground station transmitters. (As mentioned previously, this arrangement has been arrived at through some negotiation for our experiment, but could prove more difficult for very remote observatory locations.) The ACTS control personnel at NASA/Lewis in Cleveland then connect the satellite with each HDR, and verify that the signal strength is sufficient. Barring hardware complications and rain at either HDR location, the network is generally available within a few minutes of the scheduled time.

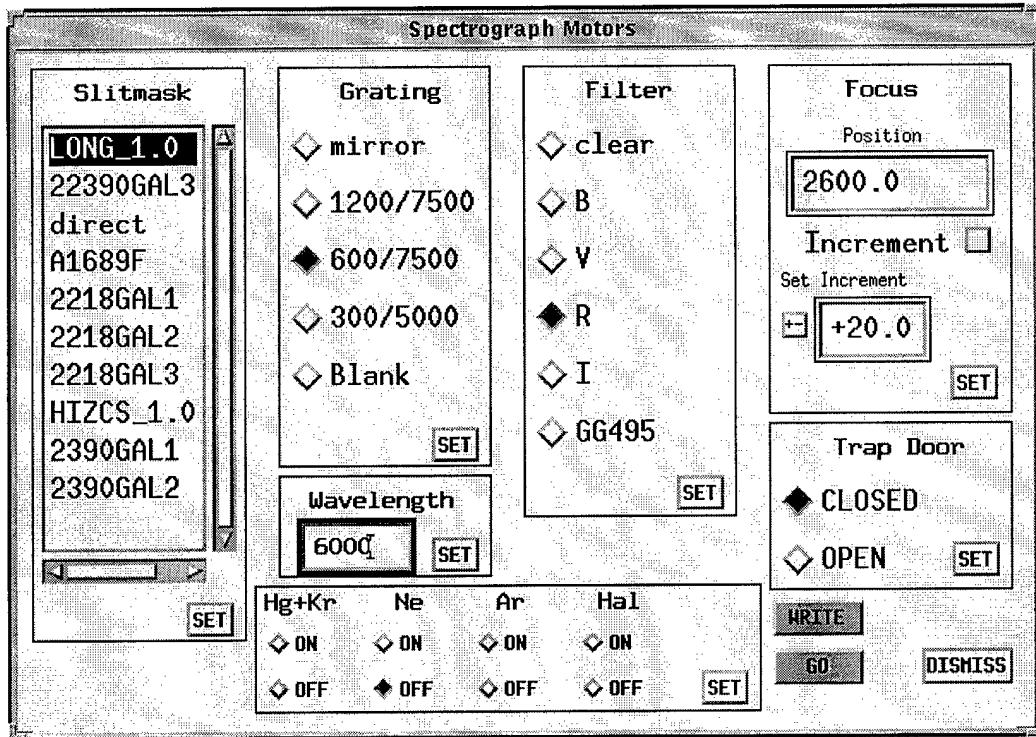


Figure 13: The LRIS instrument control interface at Keck Observatory. This window allows the user to modify the configuration of the instrument.

Contact Keck Once the network has been established, the observer customarily contacts the observing assistant (OA) at the telescope to indicate that they are ready, to check the weather at the telescope, etc. Also at this stage, an audio link is established over the ATM network, in order to alleviate the need to manipulate (and pay for!) an all-night phone call between the astronomers and the OA. For these purposes, we use a TCP-based audio tool called `rat`. Based on an earlier tool (`vat`), `rat` was developed for the MBone (Multicast Backbone), but contains full support for point-to-point audio connections. Any of the Internet telephony products commonly available could be used for this purpose. Many include useful features such as echo suppression and voice-activated microphones.

LRIS hardware setup In order to use the LRIS instrument remotely, control must be transferred from the normal instrument control computer at Keck to the duplicate machine which is connected to the ATM network. This involves a 5-minute procedure during which the VME crates which directly control the instrument are instructed to use the alternate machine, and are then rebooted. As described in the previous section, in principle this step is unnecessary, but in practice, security concerns, processing load distribution issues, high-speed networking complications, and the desire to not affect standard Keck observing techniques has led us to establish a separate remote observing control computer. In the future, this rather awkward step should disappear from the remote observing procedure.

LRIS software setup While the OA initializes the telescope control systems, the observer should then start the LRIS instrument control environment. Just as when observing on-site, a single command/menu option starts up all of the necessary tools, with the display optionally

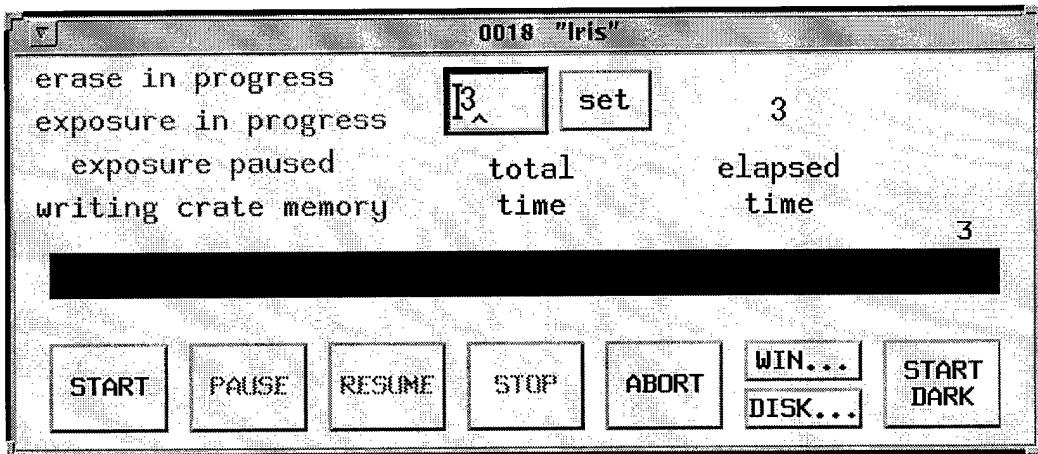


Figure 14: The LRIS instrument control interface at Keck Observatory. This window controls exposures with the CCD camera.

redirected to a remote machine. The observer can then verify any requested configuration changes, such as special slitmasks or filters, and set up personalized instrument configuration files.

LRIS operation When both the OA and observer are prepared, the observing session runs in the normal fashion: telescope moves and guiding are handled by the OA, while the observer controls the instrument configuration and CCD exposures. The LRIS instrument is controlled via a graphical interface which provides a schematic of the instrument and its current configuration (Figures 12), along with standard graphical elements (e.g., buttons, lists, etc.) for changing the configuration [3] (Figure 13). The CCD camera is controlled through a simple interface which allows the user to set and monitor exposure times (Figure 14). Details such as the number of output CCD amplifiers and the image save directory can be specified as well, of course. A real-time image display is provided, based on the FIGDISP software from FIGARO (Figure 15). Should questions arise, documentation for LRIS and its software are available to the observer via the World Wide Web³. The figure on the cover of this report illustrates one of the first images taken remotely with Keck using the ACTS high-speed network.

Should any errors be indicated by either the instrument or detector control systems, the OA and/or engineer can be called upon to examine the problem, as with on-site observing. In such situations, we often use a collaborative whiteboard tool, *wb*, which allows text and graphics to be transmitted in real-time to both parties. This is useful for indicating error messages, describing image characteristics, transmitting numerical values, etc. (As an aside, we note that *wb* and the rest of the software used for this project is available for free, with the exception of the TCP-LFN software from Sun Consulting.) Should problems arise with the network, personnel may be contacted at the ACTS control center and/or the HDR sites.

Finally, in addition to the instrument control software, all of the usual observing software tools are available remotely: telescope pointing and UT meters, guider window eavesdropping images, etc. Of course, standard TCP tools such as *telnet* and *ftp* are used regularly to retrieve images to the local system, where any of several data reduction packages com-

³<http://www2.keck.hawaii.edu:3636/>

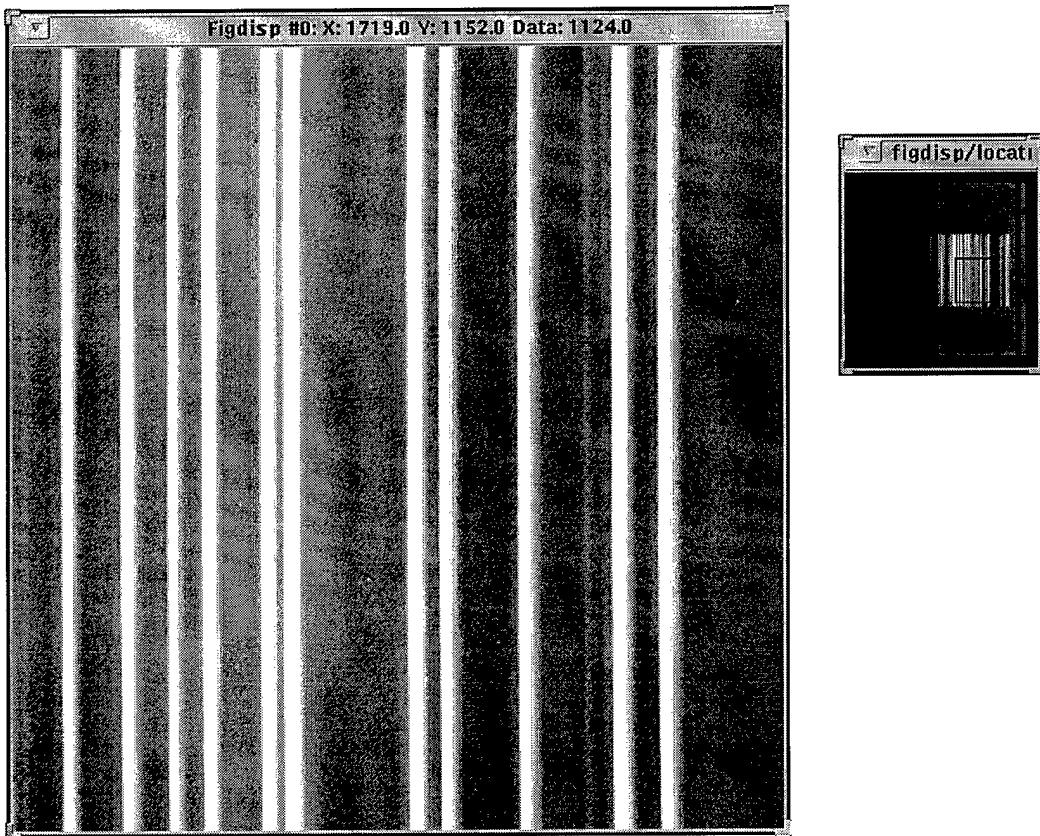


Figure 15: The LRIS instrument control interface at Keck Observatory. This window provides a real-time graphical interface for the observations. Images are transferred directly to this display as the CCD chip is read out.

monly used in astronomical image processing are available. As mentioned above, one of the outstanding features of remote observing is the wealth of familiar software and hardware facilities that are available at the user's home institution: printers, personal workstations and software, libraries, etc.

System shutdown Following the end of the observing session, control of the LRIS instrument is returned to the primary computer. This leaves the instrument ready for the next day's engineering and non-remote observing. Finally, a phone call is usually made to the ACTS control center to verify that the run is completed.

The remote observing process was developed and optimized over the lifetime of the project, with problem solution procedures being the slowest to crystallize. The final several observing runs of the project ran smoothly from a procedural standpoint, as both the astronomers and the OAs became more familiar with this mode of operation. User interface improvements, such as menu selections for commonly used remote tools, have been added to minimize the impact of the remote connection on the observer and the OA. With the possible exception of the re-boot of the instrument control crates, it is our conclusion that the average astronomer could observe remotely on Keck with at least as much ease as actually traveling to the observatory. (In either case it would be recommended that first-time users collaborate with more experienced ones, as with any new instrument or telescope.)

6 FUTURE WORK

Due to the limited scope of the ACTS project, future work from the standpoint of Keck Observatory will be concerned with establishing a more permanent remote observing facility via a ground-based network. Many of the software tools and procedures from this project will be immediately applicable to such a system. There are currently two projects underway toward this goal.

Remote observing from the Keck Headquarters in Waimea, Hawaii, has been implemented for both Keck Telescopes. The network in use between Headquarters and the observatory on the summit of Mauna Kea has slowly evolved to its current form, a T-3 (45 Mbit/sec) fiber link. This bandwidth is more than sufficient to allow remote control of the instrument, image data downloading, and eavesdropping operations. The instrument control software is networked in the same way as for our remote satellite-based system: remote display of windows using X Window System protocols.

Every attempt has been made at Headquarters to emulate on-site observing at the telescope, including separate remote control rooms for each telescope, identical software environments, and astronomer quarters. A single T-1 of bandwidth is used by an advanced PictureTel videoconferencing system, which keeps the observers and OAs in video and audio contact for the entire night. Although a primary benefit of remote observing from California is not realized, namely the reduction of travel time and costs, the proximity of technical observatory staff and the freedom from altitude-related difficulties has made remote observing from Keck Headquarters a very popular mode of observing. As much as 75% of the observing in a given month currently takes place remotely at Headquarters (depending primarily on the complexity of the instrument being used).

A separate project has been undertaken by Bob Kibrick and others at the University of California Observatories (UCO/Lick) and Keck Observatory to enable remote observing from California over terrestrial networks. Initial experiments have used the Internet, with the eventual goal to acquire the necessary bandwidth in the form of a guaranteed-bandwidth leased line. The key to this project has been the decision to remove the instrument control interface to the remote computer, with only low-level command packets and image data packets being transferred over the network. This minimizes the traffic over the network, while enabling a quickly responding interface. This separation of the user interface from the underlying instrument control software has proved relatively easy to implement because of the modular construction of the Keck Telescope control system. Figure 11 illustrates that this separation of the "observing control computer" and the "LRIS control computers" already exists, and was in fact used to our advantage in the ACTS project as well, to create a separate remote observing control computer connected to the ATM network.

The other advantage to this approach is that it enables the large image data files to be transferred to the remote host while the instrument CCD is reading the data from the hardware itself. Since this operation currently takes one or two minutes with the LRIS instrument (depending on the instrument mode), the result is that the remote user will see exactly the same image "read-out" rate as the local user, provided that the bandwidth is at least 8 Mbytes per 60 seconds, or 1.1 Mbit/sec, less than a T-1. Although this is an instrument-specific number and the required bandwidth will certainly increase as instruments become more complex with larger detectors, this allows us to implement an initial remote observing system with little software or hardware expense. This remote observing technique was successfully demonstrated at the SPIE meeting in July of last year [4]. A similar test in October has demonstrated the ability of this technique to multiplex several remote users at different locations, a capability of great benefit to large observing teams. We expect remote observing from the mainland U.S. to become increasingly common with Keck Observatory in the next few years.

7 CONCLUSIONS

The Keck/ACTS remote observing project has allowed us to experiment with remote observing on the Keck Telescope with a bandwidth not yet available to Hawaii via terrestrial networks. We have compared a variety of paradigms and tools for remote operations, and evaluated the performance of local tools over a large-bandwidth, long-delay network connection. Our work has led to a number of conclusions, many of which are applicable to ground-based remote observing efforts and non-astronomical communications satellite experiments:

1. Remote observing techniques have the potential to save appreciable expenditures in terms of money and time, while simultaneously enabling increased levels of collaboration in the observing process. In the case of an observatory with large numbers of observers, short observing runs, and/or a very remote site, these savings may easily outweigh network costs to enable remote observing.
2. The portable design of the Keck Telescope and instrument control systems has enabled remote observing to be implemented with only relatively minor software modifications. However, additional tools are needed over those available on-site to create a collaborative environment among the remote observing astronomers and the on-site telescope staff. Such tools are becoming widely available with the expansion and increasing popularity of the Internet.
3. At the current time, high-speed terrestrial networks are the most viable source for adequate bandwidth to enable true remote observing. While the ACTS system is not sufficiently robust to enable remote observing, this testbed project suggests that future commercial-grade communications satellites may provide the reliability and affordability necessary for high-bandwidth remote software applications.
4. The most outstanding problem regarding the viability of geosynchronous communications satellites for Internet-based software applications concerns the performance of the standard TCP/IP protocol over high-bandwidth, long-delay time networks. Although the initial set of extensions (i.e., RFC 1323) provide some relief, and several groups (e.g., Mitre Corporation) are working on this problem, its solution may determine the ultimate role for satellite communications in the WAN market.

8 ACKNOWLEDGMENTS

The research described in this paper was performed primarily at the California Institute of Technology and at the Center for Space Microelectronics Technology at the Jet Propulsion Laboratory. It has been sponsored by the National Aeronautics and Space Administration, Office of Advanced Concepts and Technology, under grant BK-509-20-42-00-00. The W. M. Keck Observatory is operated as a scientific partnership between the California Institute of Technology and the University of California. It was made possible by the generous gift of the W. M. Keck Foundation, and the support of its late president, Howard Keck. We also wish to thank Jimi Patel and Carl McFadden, our ground station operators at JPL and Tripler Army Medical Center, respectively, and Jon Chock, the Keck Observatory Systems Administrator, for their crucial assistance for the entire duration of this project. The scheduling and operations staff for ACTS have been extremely accommodating as well.

9 BIBLIOGRAPHY

- [1] J. D. Cavanaugh and T. J. Salo. Internetworking with ATM WANs. December 14, 1992.
- [2] N. K. Cheung. The infrastructure for gigabit computer networks. *IEEE Communications Magazine*, pages 60–68, April 1992.
- [3] J. G. Cohen, J. L. Cromer, S. Southard Jr., and D. Clowe. The software for the LRIS on the Keck 10-meter telescope. In D. R. Crabtree, R. J. Hanisch, and J. Barnes, editors, *Astronomical Data Analysis Software and Systems III*, volume 61, page 469, 1994.
- [4] A. Conrad, J. Gathright, and R. I. Kibrick. Remote observing with the Keck Telescopes. In H. Lewis, editor, *Telescope Control Systems II*, volume 3112, 1997.
- [5] V. Jacobson, R. Braden, and D. Borman. *RFC 1323: TCP Extensions for High Performance*. May, 1992.
- [6] J. Nelson. The Keck Telescope. *Am. Sci.*, 77:170, 1989.
- [7] J. B. Oke, J. G. Cohen, M. Carr, J. Cromer, A. Dingizian, F. H. Harris, S. Labrecque, R. Lucinio, W. Schaal, H. Epps, and J. Miller. The Keck Low Resolution Imaging Spectrometer. *Pub. Ast. Soc. Pac.*, 107:375, 1995.
- [8] J. W. Percival. Remote observing from the bottom up: the architecture of the WIYN telescope control system. In P. T. Wallace, editor, *Telescope Control Systems*, volume 2479, pages 33–40, 1995.
- [9] S. Scott and R. Finch. The new user interface for the ovro millimeter array. In R. Albrecht, editor, *Astronomical Data Analysis Software and Systems VII*, 1998.
- [10] R. J. Vetter and D. H. C. Du. Issues and challenges in ATM networks. *Communications of the ACM*, 38:28–109, 1995.
- [11] D. G. York. The 3.5 meter at Apache Point Observatory. *Bull. Am. Ast. Soc.*, 186:44.04, 1995.
- [12] A. A. Zijlstra, A. Wallander, L. Kaper, and J. A. Rodriguez. Remote observing at the ESO NTT and CAT telescopes. *Pub. Ast. Soc. Pac.*, 109:1256, 1997.